

Summary

Improving our understanding of hurricane inter-annual variability and the impact of climate change (e.g., doubling CO₂ and/or global warming) on hurricanes brings both scientific and computational challenges to researchers. As hurricane dynamics involves multiscale interactions among synoptic-scale flows, mesoscale vortices, and small-scale cloud motions, an ideal numerical model suitable for hurricane studies should demonstrate its capabilities in simulating these interactions. The newly-developed multiscale modeling framework (MMF, Tao et al., 2007) and the substantial computing power by the NASA Columbia supercomputer show promise in pursuing the related studies, as the MMF inherits the advantages of two NASA state-of-the-art modeling components: the GEOS4/fvGCM and 2D GCEs. This article focuses on the computational issues and proposes a revised methodology to improve the MMF's performance and scalability. It is shown that this prototype implementation enables 12-fold performance improvements with 364 CPUs, thereby making it more feasible to study hurricane climate.

Improvements in the Scalability of the NASA Goddard Multiscale Modeling
Framework for Hurricane Climate Studies

Bo-Wen Shen^{1,2}, Wei-Kuo Tao¹, and Jiun-Dar Chern^{1,3}

¹Laboratory for Atmospheres
NASA/Goddard Space Flight Center
Greenbelt, MD 20771, USA

²Earth System Science Interdisciplinary Center
University of Maryland, College Park
College Park, MD 20742, USA

³Goddard Earth Science and Technology Center
University of Maryland, Baltimore Country
Baltimore, MD 21228

March 2007

Submitted to the
International Journal of High Performance Computing Application

Short title: Improving the MMF's Scalability

Dr. Bo-Wen Shen
UMCP/ESSIC
Laboratory for Atmospheres, Code 613
NASA/Goddard Space Flight Center
Greenbelt, MD 20771, USA
Phone: (301) 614-6251
bo-wen.shen.1@gsfc.nasa.gov

Dr. Wei-Kuo Tao
Mesoscale Atmospheric Processes Branch, Code 613.1
Laboratory for Atmospheres
NASA/Goddard Space Flight Center
Greenbelt, MD 20771, USA
Phone: (301) 614-6269
wei-kuo.tao.1@gsfc.nasa.gov

Dr. Jiun-Dar Chern
UMBC/GEST
Mesoscale Atmospheric Processes Branch, Code 613.1
Laboratory for Atmospheres
NASA/Goddard Space Flight Center
Greenbelt, MD 20771, USA
Phone: (301) 614-6175
jiun-dar.chern.1@gsfc.nasa.gov

Summary:

A current, challenging topic in hurricane research is how to improve our understanding of hurricane inter-annual variability and the impact of climate change on hurricanes. Paired with the substantial computing power of the NASA Columbia supercomputer, the newly-developed multi-scale modeling framework (MMF, Tao et al., 2007) shows potential for the related studies. The MMF consists of two NASA state-of-the-art modeling components, including the finite-volume General Circulation Model (fvGCM, Lin et al., 2004) and the Goddard Cumulus Ensemble model (GCE, Tao et al., 1993, 2003). For hurricane climate studies, the MMF's computational issues need to be addressed. After introducing a meta grid system, we integrate the GCEs into a meta-global GCE in this grid-point space, and apply a 2D domain decomposition. A prototype parallelism implementation shows very promising scalability, giving a super-linear speedup as the number of CPUs is increased from 30 to 364. This scalability improvement makes it more feasible to study hurricane climate.

1. Introduction

Studies in hurricane inter-annual variability and the impact of climate change (e.g., global warming) on hurricanes have received increasing attention (Kerr, 2006), particularly due to the fact that 2004 and 2005 were the most active hurricane seasons in the Atlantic while 2006 was not as active as predicted. Thanks to recent advancements in numerical models and supercomputer technology, these topics can be addressed better than ever before.

Earth (atmospheric) modeling activities have been conventionally divided into three major categories based on scale separations: synoptic-scale, meso-scale, and cloud (micro)-scale. Historically, partly due to limited access to computing resources, hurricane climate has been studied mainly with general circulation models (GCMs) (Bengtsson et al., 2006 and references therein) and occasionally with regional mesoscale models (MMs). The former have the advantage of simulating global large-scale flow, while the latter make it possible to simulate realistic hurricane intensity and structure with fine grid spacing. However for hurricane climate studies, the resolutions used in GCMs and MMs were still too coarse to resolve small-scale convective motion, and therefore “cumulus parameterizations” (CPs) were required to emulate the effects of unresolved subgrid-scale motion. For example, a CP and cumulus momentum transport parameterization were still applied in the high-resolution hurricane simulations by Oouchi et al. (2006), who studied tropical cyclone climatology by running a global model at a resolution of 20km on the Japan Earth Simulator. Because the development of CPs has been slow, their performance is a major limiting factor in hurricane simulations.

Though hurricane formation and intensification mechanisms are still not fully understood, it is widely accepted that “cooperative” as opposed to “competitive” interaction between large-scale flow and cloud-scale convection leads to hurricane intensification. Therefore, accurate representation of non-hydrostatic cloud-scale convection and its interaction with environmental flows is crucial in hurricane studies. Cloud-resolving models (CRMs) have been extensively developed to achieve this, aimed

at advancing the development of CPs. For example, in the Global Energy and Water Cycle Experiment (GEWEX), CRMs were chosen as the primary approach to improve the representation of moist processes in large-scale models (Randall et al., 2003a). However, all CRMs, with only one exception¹, are still executed in limited areas, making difficult to understand hurricane statistics at large temporal and spatial scales.

During the last several years, an innovative approach that applies a massive number of CRMs in a global environment has been proposed and used to overcome the CP deadlock in GCMs (Randall et al., 2003b; Tao et al., 2007). This approach is called the multiscale modeling framework (MMF) or super-parameterization, wherein a CRM is used to replace the conventional CP at each grid point of a GCM. Therefore, the MMF has the combined advantages of the global coverage of a GCM and the sophisticated microphysical processes of a CRM and can be viewed as an alternative to a global CRM. Currently, two MMFs with different GCMs and CRMs have been successfully developed at Colorado State University (CSU) and NASA Goddard Space Flight Center (GSFC), and both have produced encouraging results in terms of a positive impact on simulations of large-scale flows via the feedback of resolved convection by CRMs. Among them is the improved simulation of the Madden-Julian Oscillation (MJO, Tao et al., 2007), which could potentially improve long-term forecasts of tropical cyclones through deep convective feedback. However, this approach poses a great computational challenge for performing multi-decadal runs to study hurricane climate, because nearly 10,000 copies of the CRM need to run concurrently. These tremendous computing requirements and the limited scalability in the current Goddard MMF restrict the GCM's resolution to about 2 degree, which is too coarse to capture realistic hurricane structure. In this report, computational issues and a revised model coupling approach will be addressed with the aim of improving the Goddard MMF's capabilities for hurricane climate studies.

2. The Goddard MMF on the NASA Columbia Supercomputer

¹ The first global cloud-resolving model is being running at the Japan Earth Simulator Center (e.g., Tomita et al., 2005). However, it is still challenging to study hurricane climate with this model from both scientific and computational perspectives (e.g., Miura et al., 2007).

In late 2004, the Columbia Supercomputer (Biswas et al., 2007) came into operation with a theoretical peak performance of 60 TFLOPs (trillion floating-point operations per second) at the NASA Ames Research Center (ARC). It consists of twenty 512-cpu nodes, which give 10,240 CPUs and 20 tera-bytes (TB) of memory. Columbia achieved a performance of 51.9 TFLOPs with the LINPACK (Linear Algebra PACKage) benchmark and was ranked second on the TOP500 list in late 2004; it was still ranked at No. 8 in late 2006. The cc-NUMA (cache-coherence non-uniform memory access) architecture supports up to 1 TB shared memory per node. Nodes are connected via a high-speed InfiniBand interconnect, and each node can be operating independently. These unique features enable complex problems to be resolved with large-scale modeling systems.

The Goddard MMF is based on the NASA Goddard finite-volume GCM (fvGCM) and the Goddard Cumulus Ensemble model (GCE). While the fvGCM has shown remarkable capabilities in simulating large-scale flows and thus hurricane tracks (Atlas et al., 2005; Shen et al., 2006a,b,c), the GCE is well known for its superior performance in representing small cloud-scale motions and has been used to produce more than 90 refereed journal papers (e.g., Lang et al., 2003; Tao et al., 2003). The fvGCM is running at a $2^{\circ}\times 2.5^{\circ}$ resolution, and 13,104 GCEs are “embedded” in the fvGCM to allow explicit simulation of cloud processes in a global environment. Currently, only thermodynamic feedback between the fvGCM and the GCEs is implemented. The time step for the individual 2D GCE is ten seconds, and the fvGCM-GCE coupling interval is one hour at this resolution. Under this configuration, 95% or more of the total wall-time for running the MMF is spent on the GCEs. Thus, wall-time could be significantly reduced by efficiently distributing the large number of GCEs over a massive number of processors on a supercomputer.

Over the past few years, an SPMD (single program multiple data) parallelism has been implemented in both the fvGCM and GCE with good parallel efficiency separately (Putman et al., 2005; Juang et al., 2007). Therefore, in addition to the massive number of GCEs that need to be coupled, different parallelisms in these two models make coupling

very challenging. In the following sections, both the GCE and fvGCM are introduced as well as a revised strategy for coupling these two model components.

2.1 The Goddard Cumulus Ensemble model (GCE)

Over the last two decades, the Goddard Cumulus Ensemble model (GCE) has been developed in the mesoscale dynamics and modeling group, led by Dr. W.-K. Tao, at NASA Goddard Space Flight Center. The GCE has been well tested and continuously improved. The model's main features were described in detail in Tao and Simpson (1993) and Tao et al. (1993), and its recent improvements were documented in Lang et al. (2003) and Tao et al. (2003). Table 1 gives a summary of the major characteristics of the GCE. Typical model runtime configurations are (a) (256, 256) grid points in the (x, y) directions with a grid spacing of 1-2 km; (b) 40-60 vertical stretched levels with a model top at 10-50 hPa; (c) open or cyclic lateral boundary conditions; and (d) a time step of 6 or 12 seconds. Fig. 1 shows a cloud visualization from a high-resolution simulation.

The GCE has been implemented with a 2D domain decomposition using MPI-1 (Message Passing Interface version 1) to take advantage of recent advances in supercomputing power (Juang et al., 2007). To minimize the changes in the GCE, implementation was done with a separate layer added for data communication, which preserves all of the original array indices. Therefore, not only code readability for existing modelers/users but also code portability for computational researchers is maintained. In addition to "efficiency" enhancement, tremendous efforts were made to ensure reproducibility in simulations with different CPU layouts. Without this, it would be difficult for model developers to test the model with new changes and to compare long-term simulations generated with different numbers of CPUs.

The scalability and parallel efficiency of the GCE's parallelism implementation was extensively tested on three different supercomputing platforms: an HP/Compaq (HALEM), an IBM-SP Power4, and an SGI Origin 2000 (CHAPMAN). For both anelastic and compressible versions of the GCE, 99% parallel efficiency can be reached

with up to 256 CPUs on all of the above machines (Fig. 2). Recently, the 3D version of the GCE was ported onto the NASA Columbia supercomputer, and an attempt to scale the model beyond one 512-cpu node is being made, which can be used to help understand the applicability of running massive numbers of 3D GCEs in the MMF environment.

2.2 The finite-volume General Circulation Model (fvGCM)

Resulting from a development effort of more than ten years, the finite-volume General Circulation Model (fvGCM) is a unified numerical weather prediction (NWP) and climate model that can run on daily, monthly, decadal, or century time-scales. It has the following major components: (1) finite-volume dynamics (Lin, 2004), (2) physics packages from the NCAR Community Climate Model Version 3 (CCM3, Kiehl et al., 1996), and (3) the NCAR Community Land Model Version 2 (CLM2, Dai et al., 2003). The model was originally designed for climate studies at a coarse resolution of about 2x2.5 degree in the 1990s, and its resolution was increased to 1 degree in 2000 and 1/2 degree in 2002 for NWP (e.g., Lin et al., 2003, 2004).

The parallelization of the fvGCM was carefully designed to achieve efficiency, scalability, flexibility, and portability. Its implementation had a distributed- and shared-memory two-level parallelism², including a coarse grained parallelism with MPI (MPI-1, MPI-2, MLP, or SHMEM) and fine grained parallelism with OpenMP (Putman et al., 2005). The model's dynamics, which require a lot of inter-processor communication, have 1D MPI/MLP/SHMEM domain decomposition in the y direction and OpenMP multithreading in the z direction. One of the prominent features in the implementation is to allow multi-threaded data communication. The physical part was parallelized with the 1D domain decomposition in the y direction inherited from the dynamics part and further

² During the early stages of the parallelization of the fvGCM, the multiple-level parallelism (MLP) with a collection of Unix *fork/mmap* functions (Taft 2001) was first implemented for data communication. Thus, the two-level parallelism indeed becomes shared- and share-memory parallelism. Later, asynchronous two-sided communication with MPI-1 and one-sided communication with either MPI-2 or SHMEM were implemented. To simplify discussion in this article, the term “MPI” used along with the fvGCM will be referred to as any one of these communication paradigms.

enhanced with an OpenMP loop-level parallelism in the decomposed latitudes. CLM2 was also implemented with both MPI and OpenMP parallelism, allowing its grid cells to be distributed among processors. Between the dynamical grid cells and land patches, a data mapping (or redistribution) is required.

The fvGCM can be executed either in a serial, pure MPI, pure OpenMP, or MPI-OpenMP hybrid mode, and has been ported and tested across a variety of platforms (e.g., IBM SP3, SGI O3K, SGI Altix, Linux boxes, etc) with different Fortran compilers (e.g., Intel, SGI, IBM, DEC ALPHA, PGI, Lahey, etc). Bit-by-bit reproducibility is ensured on the same platform with different CPU layouts and/or different communication protocols. All of these capabilities speedup model development and tests, thereby making the model very robust. Fig. 3 shows the model's performance and scalability based on benchmarks with 7-day NWP runs at a 0.5° resolution³ on three different platforms: Columbia (SGI Altix 4700), Halem (DEC ALPHA), and Daley (SGI O3K). Remarkable scalability was obtained with up to about 250 CPUs. In terms of throughput, the fvGCM could simulate 1110 model days (3+ years) per wall-clock day (days/day) with 240 CPUs on Columbia, 521 days/day with 288 CPUs on Halem, and 308 days/day with 300 CPUs on Daley. Even though these results are not listed for direct comparison due to different interconnect and CPU technologies (e.g., different CPU's clock speeds and cache sizes, etc), it should be noted that a 20% performance increase on Columbia is obtained with the recent upgrades (e.g., an upgrade to the Altix 4700 from the Altix 3000).

2.3 Application of the fvGCM to hurricane forecasts

After being substantially improved and tested, the fvGCM at 0.5° resolution was first run in a weather mode experimentally in early 2002 (e.g., Lin et al., 2003). As Columbia was being built in early 2004, a higher resolution (0.25°) fvGCM was deployed to perform quasi-realtime hurricane forecasts. Though hurricane prediction poses a

³ A resolution of $2 \times 2.5^\circ$ is being used in the fvGCM within the MMF, and 1° is the target resolution in this study. Thus, 0.5° should be sufficient for now. Benchmarks at higher resolution (e.g., 0.25°) are being performed on Columbia and will be documented in a separate study.

challenge for GCMs because of insufficient horizontal resolution, the 0.25° fvGCM, which doubles the resolution adopted by major weather centers at that time, was one of the first GCMs that could produce realistic tropical weather systems and remarkable hurricane forecasts (e.g., Atlas et al., 2005). While doubling the resolution of a global (NWP) model requires an 8-16X increase in computational power, the unprecedented computing capacities afforded by the NASA Columbia supercomputer allowed for a rapid increase in resolution of the fvGCM to 0.125° in early 2005 (e.g., Shen et al., 2006a) and to 0.08° in the middle of 2005, making it as one of a few global mesoscale models. As shown in Fig. 4, the first global 5-day forecast of total precipitable water with the 0.125° fvGCM produced not only an accurate forecast of the large-scale flow but also very realistic fine structures for tropical systems, including the landfalling hurricane Frances (2004).

The 2005 Atlantic hurricane season was the most active in recorded history. There were twenty-eight tropical storms and fifteen hurricanes, four of which were Category 5 hurricanes. Though hurricane track forecasts have been steadily improved, progress on intensity forecasts has been very slow over the past decades. The performance of the $1/8$ degree fvGCM on hurricane intensity predictions was first demonstrated with six 5-day forecasts of Hurricane Katrina, showing remarkable forecasts with errors in central pressure of only ± 12 hPa (Shen et al., 2006b). Accurate 5-day track forecasts and realistic vertical structures for Katrina are shown in Figs. 5 and 6, respectively. A systematic study on the model's ability to accurately predict the track and intensity of intense hurricanes in 2004 and 2005 is being conducted; preliminary analyses have been documented in Shen et al. (2006c). It was found that the global mesoscale model shows promise in improving short-term hurricane predictions. However, for hurricane climate studies, long-term simulations with enabled or disabled CPs still hold uncertainties. In comparison, the MMF with the combined advantages of the fvGCM and GCE might provide an alternative solution, if its capabilities can be extended as discussed below.

3. Results and Discussion on the Enhanced MMF

The Goddard MMF implementation consists of the fvGCM at $2^\circ \times 2.5^\circ$ resolution and 13,104 GCEs, each of which is embedded in one grid cell of the fvGCM (Fig. 7). Since it would require a tremendous effort to implement an OpenMP parallelism into the GCE or to extend the 1D domain decomposition to 2D in the fvGCM, the MMF only inherited the fvGCM’s 1D MPI parallelism, though the fvGCM was parallelized with both MPI and OpenMP paradigms. This single-component approach limited the MMF’s scalability to 30 CPUs, and thereby posed a challenge for increasing the resolution of the fvGCM and/or extending the GCE’s dimension from 2D to 3D. To overcome this difficulty, a different strategic approach is needed to couple the fvGCM and GCEs.

From a computational perspective, the concept of “embedded GCEs” should be completely forgotten, as it restricts the view on the data parallelism of the fvGCM. Instead, the 13,104 GCEs should be viewed as a *meta global GCE* (mgGCE) in a *meta gridpoint system*, which includes 13,104 grid points. This grid system, which is not tied to any specific grid system, is assumed to be the same as the latitude-longitude grid structure in the fvGCM for convenience. With this concept in mind, each of the two distinct parts (the fvGCM and mgGCE) in the MMF could have its own scaling properties (Fig. 8). Since most of wall-time was spent on the GCEs, the wall-time could be substantially reduced by deploying a highly scalable mgGCE and/or coupling the mgGCE with the fvGCM using an MPMD (multiple programs multiple data) parallelism.

Data parallelism in the mgGCE indeed becomes a task parallelism, namely distributing 13,104 GCEs among processors. Because cyclic lateral boundary conditions are used in each GCE, the mgGCE has no ghost region in the meta grid system and can be scaled “embarrassingly” with a 2D domain decomposition. For the coupled MMF, which has major overhead only in data redistribution (or data regridding) between the fvGCM and the mgGCE, its scalability and performance will depend mainly on the scalability and performance of the mgGCE and the coupler, which is the interface between the fvGCM and mgGCE. Under this current definition, a grid inside each GCE, running at one meta grid, becomes a *child grid* (or sub-grid) with respect to the parent (meta) grid (Fig. 8). Since an individual GCE can still be executed with its native 2D MPI

implementation in the child grid-point space, this second level of parallelism can greatly expand the number of CPUs. Potentially, the coupled MMF along with the mgGCE could be scaled at a multiple of 13,104 CPUs. Having two different components, this coupled system is also termed a multi-scale multi-component modeling framework in this study.

Another advantage of introducing the mgGCE component is to allow the adoption of the idea of land-sea masks used in a land model. For example, if computing resources are limited, a cloud-mask file can be used to specify limited regions where the GCEs should be running. A more sophisticated cloud-mask implementation in the mgGCE will enable one to choose a variety of GCEs (2D vs. 3D, bulk vs. bin microphysics) depending on geographic location. Thus, computational load balances can be managed efficiently.

To achieve all of the aforementioned functionalities, a scalable and flexible coupler and a scalable parallel I/O module need to be developed. The coupler should be designed carefully in order: (1) to minimize the changes in the GCE and permit it as a stand-alone application or a one element/component in the mgGCE; (2) to seamlessly couple the mgGCE and fvGCM to allow for a different CPU layout in each of these components; (3) to allow the mgGCE to be executed in a global, channel, or regional environment with a suitable configuration in the cloud-mask file. A scalable (parallel) I/O module needs to be implemented in the meta grid-point space, since it is impractical to have the individual GCE to do its I/O.

As a stand-alone model, the mgGCE can be also tested offline with large-scale forcing derived from model reanalysis [e.g., from the Global Forecast System (GFS) at the National Centers for Environmental Prediction (NCEP)] or from high-resolution model forecasts (e.g., the fvGCM). To assure the implementation in the mgGCE is correct, simulations with the mgGCE at a single meta point should be identical to those with a regular GCE. One potential application of the mgGCE is to investigate the short-term evolution of hurricane Katrina's (2005) precipitation by performing simulations driven by the NCEP GFS T382 (~35km) reanalysis data at a 6h time interval. This approach can be further extended by replacing the GFS reanalysis by 1/8° fvGCM

forecasts at a smaller time interval (see more detailed information about these forecasts in Shen et al., 2006b).

At this time, a prototype MMF including the mgGCE, fvGCM and coupler has been successfully implemented. The technical approaches are briefly summarized as follows: (1) a master process allocates a shared memory arena for data redistribution between the fvGCM and mgGCE by calling the Unix *mmap* function; (2) the master process spawns multiple (parent) processes with a 1D domain decomposition in the y direction by a series of Unix *fork* system calls; (3) each of these parent processes then forks several child processes with another 1D domain decomposition along the x direction; (4) data gathering in the mgGCE is done along the x direction and then the y direction; (5) synchronization is implemented with the atomic *_sync_add_and_fetch* function call on the Columbia supercomputer. While steps (1), (2), and (5) were previously used in MLP (multiple level parallelism) by Taft (2001), this methodology is now extended to the multi-component system.

Fig. 9 shows preliminary benchmarks with very promising scalability up to 364 CPUs. Here the speedup is determined by T_{30}/T , where T is the wall time to perform a 5-day forecast with the MMF and T_{30} the time spent using 30 CPUs. The run with 30 CPUs was chosen as a baseline simply because this configuration was previously used for production runs. A speedup of (3.93, 7.28, and 12.43) is obtained for (91, 182, and 364) CPUs, respectively. As the baseline has load imbalance and excessive memory usage in the master process, it is not too surprising to obtain a super-linear speedup. Further analysis of the MMF's throughput indicates that it takes about 164 minutes to finish a 5-day forecast using 364 CPUs, which meets the requirement for performing realtime numerical weather prediction. A yearly simulation would only take 8 days to run with 364 CPUs as opposed to 96 days with 30 CPUs. This makes it far more feasible for studying hurricane climate.

4. Concluding Remarks

Improving our understanding of hurricane inter-annual variability and the impact of climate change (e.g., doubling CO₂ and/or global warming) on hurricanes brings both scientific and computational challenges to researchers. As hurricane dynamics involves multiscale interactions among synoptic-scale flows, mesoscale vortices, and small-scale cloud motions, an ideal numerical model suitable for hurricane studies should demonstrate its capabilities in simulating these interactions. The newly-developed multi-scale modeling framework (MMF, Tao et al., 2007) and the substantial computing power by the NASA Columbia supercomputer show promise in pursuing the related studies, as the MMF inherits the advantages of two NASA state-of-the-art modeling components: the fvGCM and 2D GCEs. This article focuses on the computational issues and proposes a revised methodology to improve the MMF's performance and scalability. It has been shown that this prototype implementation can improve the MMF's scalability substantially without the need to make major changes in the fvGCM and GCEs.

To achieve these goals, the concept of a meta grid system was introduced, grouping a large number of GCEs into a new component called the mgGCE. This permits a component-based programming paradigm to be used to couple the fvGCM and mgGCE. A prototype MMF is then implemented for data redistribution between these two components. This revised coupled system is also termed a multiscale multicomponent modeling framework as both the fvGCM and mgGCE are separate components with their own parallelism. This proof-of-concept approach lays the groundwork for a more sophisticated modeling framework and coupler to solve unprecedentedly complex problems with advanced computing power. For example, the cloud-mask idea associated with the mgGCE will enable GCEs to run with a variety of choices, including different dimensions (2D vs. 3D) and different microphysical packages (e.g., bulk or bin). The next step is to conduct hurricane climate studies by performing long-term MMF simulations with a channel mgGCE and 1°x1.25° fvGCM. A global channel ranging from 45°S to 45°N requires only 26024 3D GCEs with respect to 52128 GCEs for a whole globe and becomes more computationally affordable with current computing resources.

It is well known that a latitude-longitude grid system has issues such as efficiency/performance and convergence problems near the poles. As the meta grid system in the mgGCM is no longer bound to the fvGCM's grid system, this meta-grid concept could help avoid the performance issues by implementing a quasi-uniform grid system (such as a cube grid or geodesic grid) into the mgGCE. Such a deployment should lead to a substantial performance increase since 95% of the computing time for the MMF is spent on the mgGCE.

The fundamental communication paradigm for data redistribution in this implementation is similar to the MLP which was developed by Taft (2001) and used for parallelization in single-component models with tremendous benefits. The methodology is extended here to a multi-component modeling system, showing an alternative and easy way for coupling multiple components. Further improvements in the implementation include an adoption of a more portable communication paradigm (such as MPI-1 or MPI-2) and/or a sophisticated modeling framework. While the current implementation in process management, data communication/redistribution, and synchronization is solely done with Unix system calls, earlier experiences with the parallelism implementation in the fvGCM have proven that this can be easily extended with an MPI-2 implementation (Putman et al., 2005). A survey on existing frameworks such as ESMF or PRISM is being conducted; however, it is too early to make a final selection. First of all, no framework has demonstrated its superior scalability with a large number of model components and secondly this modeling system is so complex and "innovative" that it would take time for framework developers to include the MMF's requirements in their frameworks. Finally, as the bulk of the computing is done in the mgGCE, which has no ghost points, the next version of MMF with the mgGCE is envisioned to be a good candidate for meta- (grid-) computing just like the SETI@home project⁴. Namely computations in the mgGCE could be distributed among available personal- and super-computers connected via the Internet.

⁴ SETI stands for Search for Extraterrestrial Intelligence, which is a scientific experiment with computing available over the Internet. For more information see <http://setiathome.berkeley.edu/>

Acknowledgements: We thank Dr. D. Anderson for his support for developing the MMF under the NASA IDS and Cloud Modeling and Analysis Initiative (CMAI) program. The GCE is mainly supported by the NASA Headquarters Atmospheric Dynamics and Thermodynamics Program and Tropical Rainfall Measuring Mission. We would like to thank Mr. S. Lang for helpful comments on the manuscript. Finally, we'd like to express our sincere thank to the NASA Advanced Supercomputing Division and NASA Center for Computational Sciences divisions for their strong support and use of computing and storage resources.

ACRONYMS

- **ARC:** Ames Research Center
- **CLM:** Community Land Model
- **CPS:** Cumulus Parameterization Scheme
- **CRM:** Cloud-Resolving Model
- **CSU:** Colorado State University
- **fvGCM:** finite-volume General Ciculation Model
- **GCE:** Goddard Cloud Ensemble Model
- **GCM:** General Circulation Model
- **GFS:** Global Forecast System
- **GSFC:** Goddard Space Flight Center
- **mgGCE:** meta global GCE
- **MM:** Mesoscale Model
- **MMF:** Multi-scale Modeling Framework
- **MJO:** Madden Julian Oscillation
- **MLP:** Multiple Level Parallelism
- **MPMD:** Multiple Program Multiple Data
- **MPI:** Message Passing Interface
- **NASA:** National Aeronautics and Space Administration
- **NCAR:** National Center for Atmospheric Research
- **NCEP:** National Centers for Environmental Prediction
- **SPMD:** Single Program Multiple Data

References:

Atlas, R., O. Reale, B.-W. Shen, S.-J. Lin, J.-D. Chern, W. Putman, T. Lee, K.-S. Yeh, M. Bosilovich, and J. Radakovich, 2005: Hurricane forecasting with the high-resolution NASA finite volume general circulation model. *Geophys. Res. Lett.*, **32**, L03801, doi:10.1029/2004GL021513.

Bengtsson, L., K. I. Hodges, and M. Esch, 2006: Hurricane type vortices in a high-resolution global model: Comparison with observations and Re-Analyses. Submitted to *Tellus*.

Biswas, R., M.J. Aftosmis, C. Kiris, and B.-W. Shen, 2007: Petascale Computing: Impact on Future NASA Missions. *Petascale Computing: Architectures and Algorithms* (D. Bader, ed.), Chapman and Hall / CRC Press, to appear.

Dai, Y., X. Zeng, R. E. Dickinson, I. Baker, G. B. Bonan, M. G. Bosilovich, A. S. Denning, P. A. Dirmeyer, P. R. Houser, G. Niu, K. W. Oleson, C. Adam Schlosserl, and Z.-L. Yang, 2003: The Common Land Model. *Bull. Amer. Meteor. Soc.*, **84**, 1013-1023.

Juang, H.-M., W.-K. Tao, X. Zeng, C.-L. Shie, S. Lang, and J. Simpson, 2007: Parallelization of NASA Goddard Cloud Ensemble Model for Massively Parallel Computing. *Terrestrial, Atmospheric and Oceanic Sciences*. (in press)

Kerr, R., 2006: A Tempestuous Birth for Hurricane Climatology. *Science*, Vol **312**, 676-678.

Kiehl, J. T., J. Hack, G. Bonan, B. Boville, B. Briegleb, D. Williamson, P. Rasch, 1996: Description of the NCAR Community Climate Model (CCM3). NCAR Technical Note.

Lang, S., W.-K. Tao, J. Simpson and B. Ferrier, 2003: Modeling of convective-stratiform precipitation processes: Sensitivity to partitioning methods, *J. Appl. Meteor.* **42**, 505-527.

Lin, S.-J., 2004: A "vertically Lagrangian" finite-volume dynamical core for global models. *Mon. Wea. Rev.*, **132**, 2293-2307.

Lin, S.-J., B.-W. Shen, W. P. Putman, J.-D. Chern, 2003: Application of the high-

resolution finite-volume NASA/NCAR Climate Model for Medium-Range Weather Prediction Experiments. EGS - AGU - EUG Joint Assembly, Nice, France, 6 - 11 April 2003.

Lin, S.-J., R. Atlas, K.-S. Yeh, 2004: Global weather prediction and high-end computing at NASA. *Computing in Science and Engineering*, **6**(1), 29-35.

Miura, H., M. Satoh, H. Tomita, A. T. Noda, T. Nasuno, and S. Iga (2007), A short-duration global cloud-resolving simulation with a realistic land and sea distribution, *Geophys. Res. Lett.*, **34**, L02804, doi:10.1029/2006GL027448.

Oouchi, K., J. Yoshimura, H. Yoshimura, R. Mizuta, S. Kusunoki, and A. Noda, 2006: Tropical Cyclone Climatology in a Global-Warming Climate as Simulated in a 20 km-Mesh Global Atmospheric Model: Frequency and Wind Intensity Analyses. *J. Meteo. Soc. Japan*, **84**, 259-276.

Putman, W., S.-J. Lin, and B.-W. Shen, 2005: Cross-Platform Performance of a Portable Communication Module and the NASA Finite Volume General Circulation Model. *International Journal of High Performance Computing Applications*. **19**: 213-223.

Randall, D. A., J. Curry, P. Duynkerke, S. K. Kruger, M. W. Moncrieff, B. Ryan, D. O. Starr, M. Miller, W. Rossow, G. Tseliudis, B. A. Wielikci, 2003a: The GEWEX cloud system study: A review from 2001. *Bull. Amer. Meteor. Soc.*, **84**, 455-469.

Randall D, M. Khairoutdinov, A. Arakawa, W. Grabowski, 2003b: Breaking the Cloud Parameterization Deadlock. *Bull. Amer. Meteor. Soc.*, 1547-1564.

Shen, B.-W., R. Atlas, J.-D. Chern, O. Reale, S.-J. Lin, T. Lee, and J. Chang, 2006a: The 0.125 degree finite-volume General Circulation Model on the NASA Columbia Supercomputer: Preliminary Simulations of Mesoscale Vortices. *Geophys. Res. Lett.*, **33**, L05801, doi:10.1029/2005GL024594.

Shen, B.-W., R. Atlas, O. Reale, S.-J. Lin, J.-D. Chern, J. Chang, C. Henze, and J.-L. Li, 2006b: Hurricane Forecasts with a Global Mesoscale-Resolving Model: Preliminary Results with Hurricane Katrina (2005). *Geophys. Res. Lett.*, **33**, L13813, doi:10.1029/2006GL026143.

Shen, B.-W., W.-K. Tao, R. Atlas, T. Lee, O. Reale, J.-D. Chern, S.-J. Lin, J. Chang, C. Henze, J.-L. Li, 2006c: Hurricane Forecasts with a Global Mesoscale Model on the NASA Columbia Supercomputer, AGU 2006 Fall Meeting, December 11-16, 2006.

This can be downloaded from
http://atmospheres.gsfc.nasa.gov/cloud_modeling/docs/2006_AGU_Fall_Poster.ppt

Taft, J. R., 2001: Achieving 60 gflop/s on the production cfd code overflow-mlp. *Parallel Computing*, 27(4):521-536.

Tao, W.-K., and J. Simpson, 1993: The Goddard Cumulus Ensemble Model. Part I: Model description. *Terrestrial, Atmospheric and Oceanic Sciences*, **4**, 19-54.

Tao, W.-K., and J. Simpson, C.-H. Sui, B. Ferrier, S. Lang, J. Scala, M.-D. Chou, and K. Pickering, 1993: Heating, moisture and water budgets of tropical and mid-latitude squall lines: Comparisons and sensitivity to longwave radiation. *J. Atmos. Sci.*, **50**, 673-690.

Tao, W.-K., C.-L. Shie, R. Johnson, S. Braun, J. Simpson, and P. E. Ciesielski, 2003: Convective Systems over South China Sea: Cloud-Resolving Model Simulations. *J. Atmos. Sci.*, **60**, 2929-2956.

Tao, W.-K., J. Chern, R. Atlas, D. Randall, X. Lin, M. Khairoutdinov, J. L. Li, D. E. Waliser, A. Hou, C. Peters-Lidard, W. Lau, and J. Simpson, 2007: Multi-scale modeling system: Development, applications and critical issues, *J. Geophys. Res.* (submitted).

Tomita, H, Miura, H., Iga, S., Nasuno, T., and Satoh,M.,, 2005: A Global Cloud-Resolving Simulation: Preliminary Results from an Aqua Planet Experiment *Geophys. Res. Lett.*, vol.32, L08805, doi:10.1029/2005GL022459.

Table 1

Parameters/Processes	GCE Model
Dynamics	Anelastic or compressible 2D (slab- and axis-symmetric) and 3D
Vertical Coordinate	$Z(p)$
Explicit Microphysical Processes	2-class water & 2-moment 4-class ice, 2- or 3-class ice Spectral-Bin Microphysics *
Implicit Convective Processes	Betts & Miller or Kain & Fritsch
Numerical Methods	Positive Definite Advection for Scalar Variables; 4-th Order for Dynamic Variables
Initialization	Initial Conditions with Forcing from Observations/Large-Scale Models
FDDA	Nudging
Radiation	k-distribution and four-stream discrete-ordinate scattering (8 bands) Explicit Cloud-radiation Interaction
Sub-Grid Diffusion	TKE (1.5 order)
Topography	$\Sigma_z(p)**$
Two-Way Interactive Nesting	Radiative-Type*
Surface Energy Budget	7-Layer Soil Model (PLACE) CLM - LIS TOGA COARE Flux Module
Parallelization	OPEN-MP and MPI

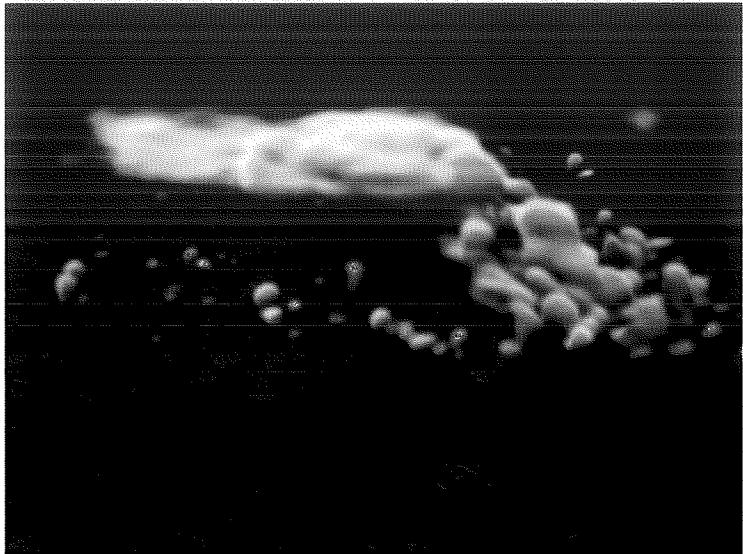


Figure 1: High-resolution simulation of the 23 Feb 1999 TRMM LBA case with the Goddard Cloud Ensemble model. Image by J. Williams of the NASA GSFC Scientific Visualization Studio.

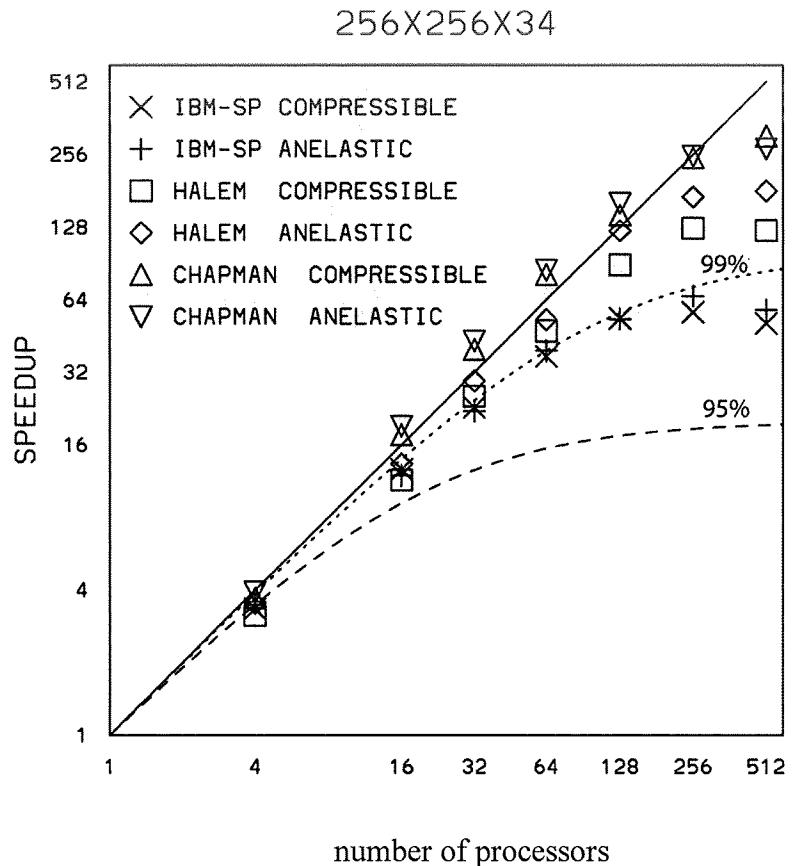


Figure 2: GCE speedup on different platforms (after Juang et al. 2007).

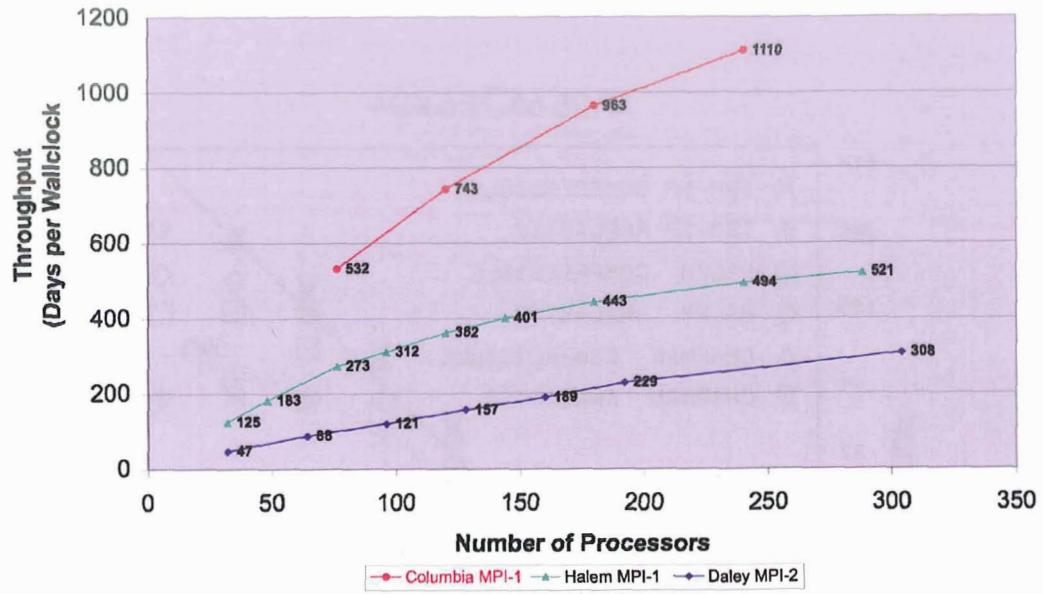


Figure 3: fvGCM's throughput (model days per wall-clock day) based on 7-day numerical weather forecasts at a $0.5^\circ \times 0.625^\circ$ resolution.

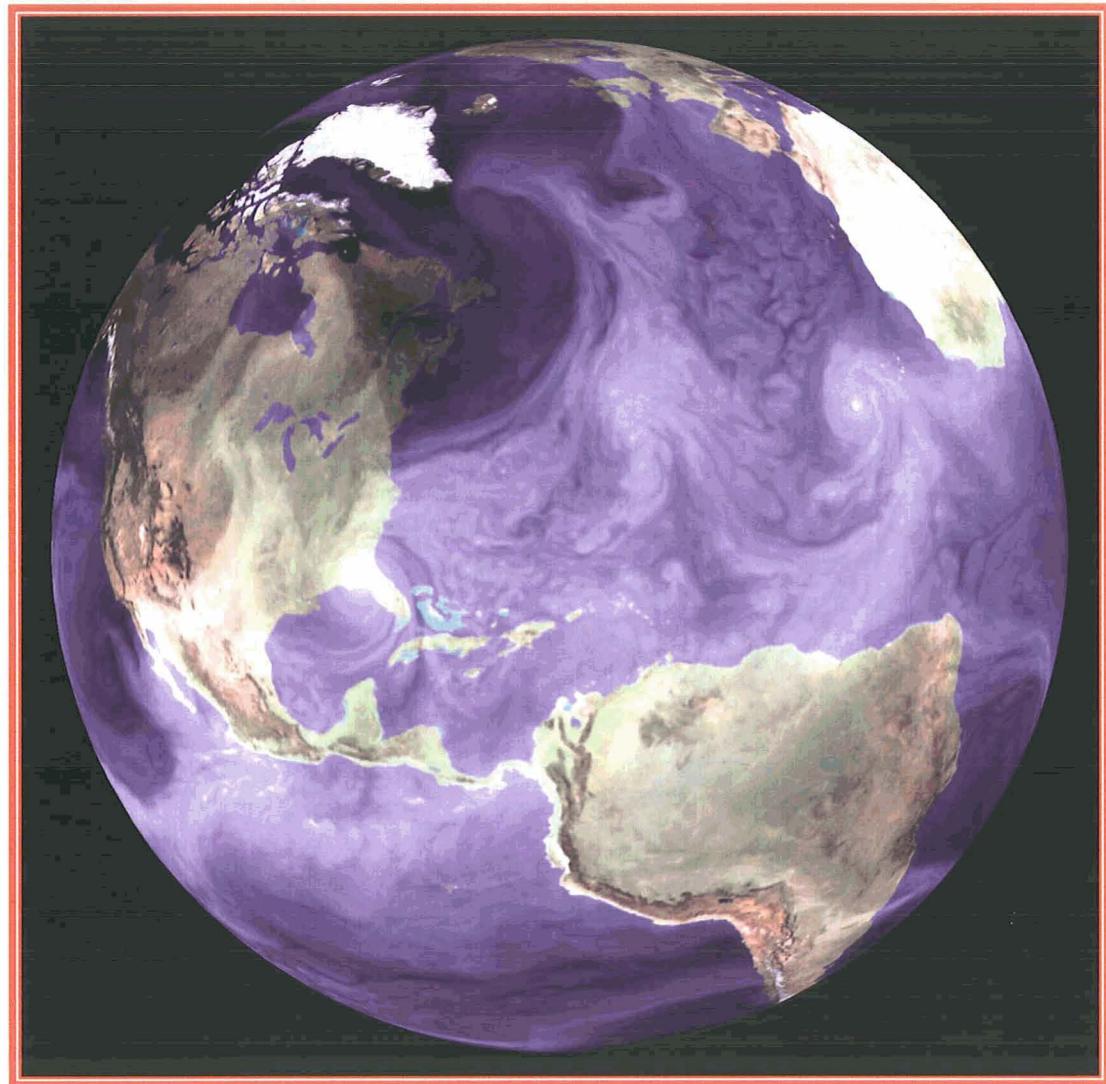


Figure 4: This global view shows total precipitable water from 5-day forecasts initialized at 0000 UTC September 1 2004 with the 1/8 degree fvGCM.

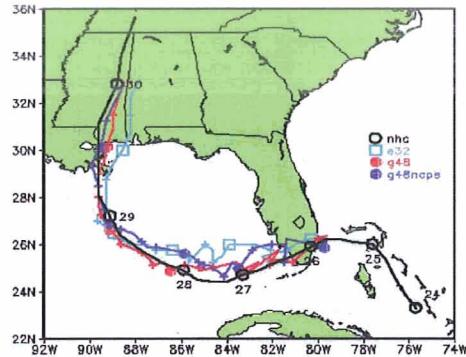


Figure 5: (a) 5-day track forecasts of hurricane Katrina (2005) initialized at 1200 UTC August 25, 2005 with the fvGCM at different resolutions: e32 (1/4 degree), g48 (1/8 degree), and g48ncps (1/8 degree without cumulus parameterizations). (courtesy American Geophysical Union, Shen et al. 2006b)

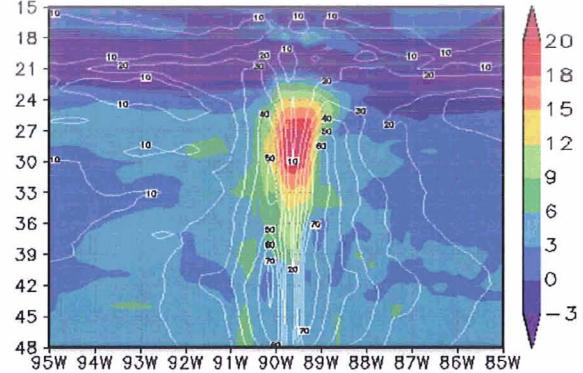


Figure 6: Simulated vertical structure of Katrina (2005) from 96h simulations with no CPs along $\text{lat}=28.5^\circ$. The vertical axis represents the model's levels. This figure shows realistic features such as horizontal maximum winds (white) near the top of the boundary layer, a narrow eyewall, and an elevated warm core (shaded).

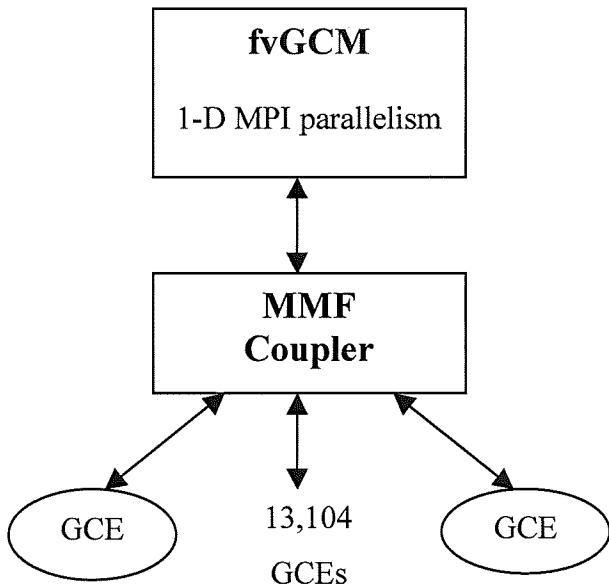


Figure 7: Schematic diagram of the fvGCM, GCEs, and MMF coupler.

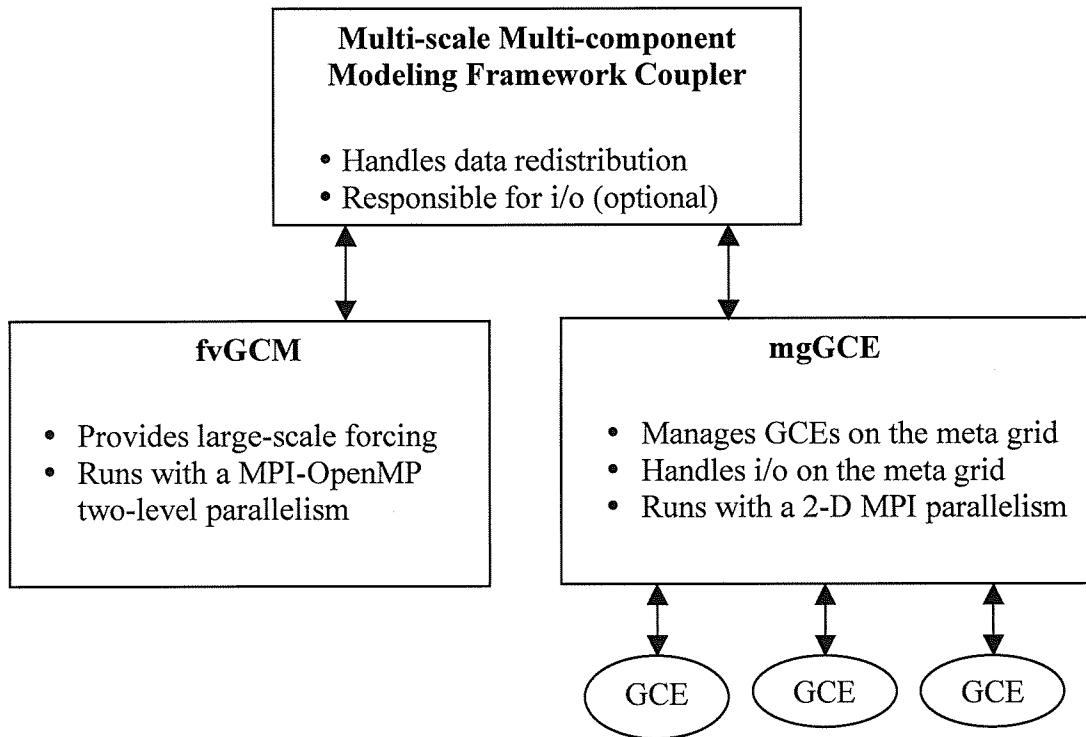


Figure 8: Schematic diagram of the meta-global GCE and a revised MMF coupler.

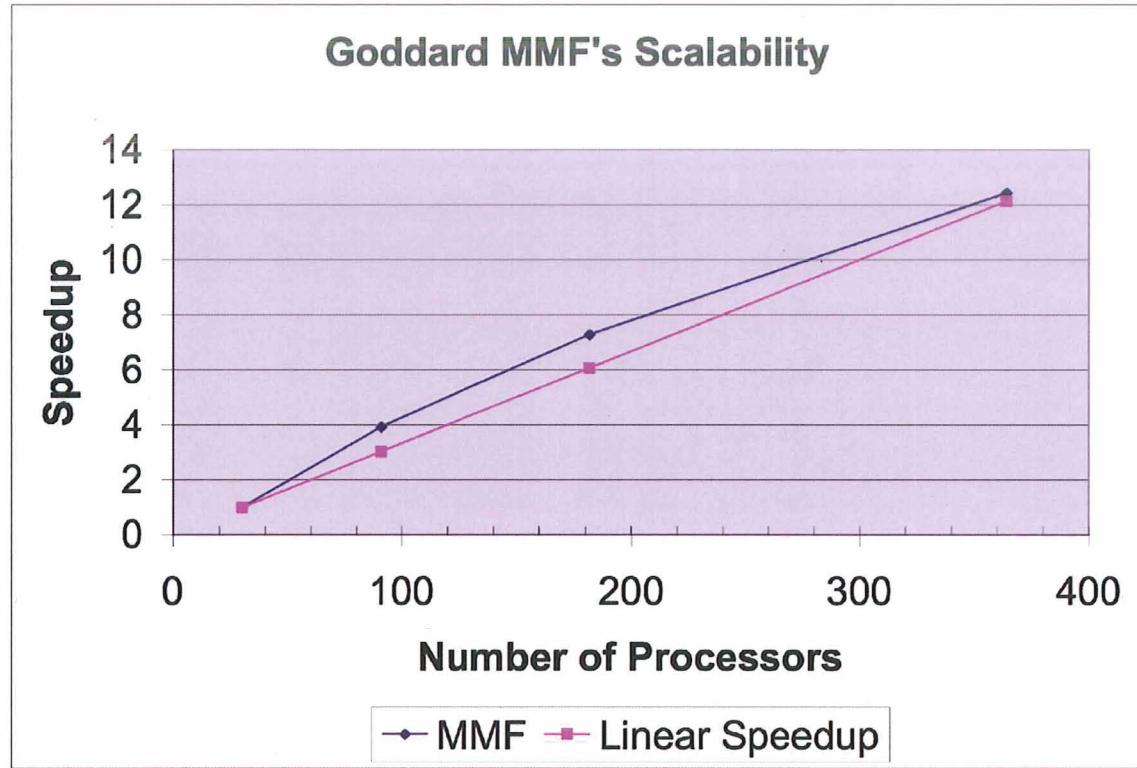


Figure 9: Scalability of the Goddard MMF with a revised parallelism.